Eurographics Symposium on Rendering (2007)
Jan Kautz and Sumanta Pattanaik (Editors)

Rapid Acquisition of Specular and Diffuse Normal Maps from Polarized Spherical Gradient Illumination

Wan-Chun Ma Tim Hawkins Pieter Peers Charles-Felix Chabert Malte Weiss Paul Debevec

University of Southern California
Institute for Creative Technologies

Abstract

We estimate surface normal maps of an object from either its diffuse or specular reflectance using four spherical gradient illumination patterns. In contrast to traditional photometric stereo, the spherical patterns allow normals to be estimated simultaneously from any number of viewpoints. We present two polarized lighting techniques that allow the diffuse and specular normal maps of an object to be measured independently. For scattering materials, we show that the specular normal maps yield the best record of detailed surface shape while the diffuse normals deviate from the true surface normal due to subsurface scattering, and that this effect is dependent on wavelength. We show several applications of this acquisition technique. First, we capture normal maps of a facial performance simultaneously from several viewing positions using time-multiplexed illumination. Second, we show that high-resolution normal maps based on the specular component can be used with structured light 3D scanning to quickly acquire high-resolution facial surface geometry using off-the-shelf digital still cameras. Finally, we present a real-time shading model that uses independently estimated normal maps for the specular and diffuse color channels to reproduce some of the perceptually important effects of subsurface scattering.

1. Introduction

Recent techniques for acquiring geometry and reflectance of objects have enabled much richer and more realistic renderings in computer graphics. Many of these techniques (e.g. [SWI, DHT*, LKG*03, GTHD, WMP*06]) use a combination of 3D scanning and photography under different lighting conditions to acquire 3D models of the object's shape and how it reflects light. When both of these characteristics are measured, the models can be used to faithfully render how the object would look from any angle, reflecting the light of any environment.

Some of the current reflectance acquisition techniques photograph an object under a small number of lighting conditions, and use this limited reflectance information to infer BRDFs across the entire object surface [LKG*03, ZERB]. However, such techniques miss some of the spatially-varying effects of specular reflections. Other techniques acquire object reflectance using many images of the object under a dense sampling of lighting directions [DHT*, CGS, WMP*06]. These techniques model spatially-varying BRDFs more accurately but are significantly more data intensive, complicating the capture of live subjects.

In this work we present a new reflectance acquisition technique that uses a small set of lighting conditions, but is able to acquire independent estimates of diffuse and specular reflectance behavior across the entire object surface. The technique uses four spherical gradient illumination patterns which effectively compute the centroid and total energy of each pixel's reflectance function, i.e. the amount of light reflected toward the camera as the object is lit from each direction on the sphere. We further show that the gradient patterns can be polarized, using either linear or circular polarization, in a way that the diffuse and specular components can be characterized independently. We show that the centroids of the diffuse and specular reflectance each yield an estimate of the pixel's surface normal, and we refer to these estimates as the *diffuse normals* and *specular normals* for the object. Thus, our four lighting patterns, each photographed under two polarization states, yield the diffuse normal (per color channel), diffuse albedo (per color channel), specular normal, and specular intensity for each pixel.

The key novel capability of our technique is estimating normal maps from the object's *specular* reflectance using only a few lighting patterns. Since few photographs are required, commonplace high-resolution still cameras can be used for the data acquisition, even for live subjects, allowing higher-resolution maps to be acquired than previous ap-

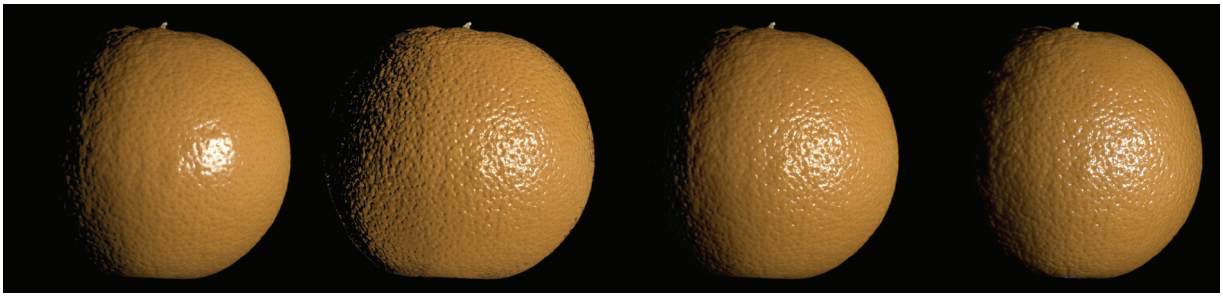

(a) Diffuse normals (b) Specular normals (c) Hybrid normals (d) Real photograph

Figure 1: *An orange rendered from acquired measurements of **(a)** diffuse normals, **(b)** specular normals, and **(c)** hybrid rendering with separate specular and diffuse normals. This data was acquired from linearly polarized spherical illumination and rendered using a manually-chosen specular roughness parameter. **(d)** A real point-lit photograph of the orange for comparison.*

proaches have allowed. Furthermore, these maps provide the best representation of high-resolution surface detail for materials exhibiting translucency. For such materials, standard 3D scanning and photometric techniques which work by analyzing only the diffuse component of the reflection will miss detail due to the blurring effects of subsurface scattering.

Our technique's normal estimates from diffuse reflection are affected by self-shadowing and mutual illumination. If left uncorrected, the diffuse normals we acquire are more useful for surface shading than for geometry measurement, akin to the *bent normals* from [Lan02] and the polynomial texture maps of [MGW]. An advantage our patterns have over traditional photometric stereo is that the lighting patterns (except under linear polarization) are view-independent, allowing normals to be estimated from any viewpoint from the same four patterns. Since our technique estimates just the normals and albedo, we use a manually-chosen specular roughness parameter for our renderings; our technique does *not* measure the full BRDF at each surface point. However, the normal maps we acquire are of a sufficiently high resolution to yield some of the appearance of spatially-varying specular roughness.

One application we show for this acquisition process leverages the fact that differences between the specular and diffuse normals are important for characterizing the reflectance of organic materials. From this observation, we present a real-time *hybrid normal* shading technique that uses independent normal maps for the specular channel and each of the three diffuse color channels. We show that this technique reproduces some of the perceptually important effects of an object's surface translucency, as seen in Fig. 1 and Fig. 10.

In another application, we show that our specular surface normal estimates can be combined with cost-effective structured light scanning to produce 3D models comparable in quality to high-resolution laser scans. Unlike laser scanning, our acquisition technique can directly acquire high-resolution surface detail of organic subjects because of its high speed and its robustness to surface translucency. We use this technique in a system which to our knowledge yields the highest-resolution photometrically-based scans of human faces to date, and has the relatively low-cost setup of two SLR cameras, a set of polarizers, and an LED sphere. In summary, our contributions include:

1. The use of a set of view-independent spherical gradient illumination patterns to recover photometric normal maps of objects based on either diffuse or specular reflectance.
2. The use of linear or circular polarization patterns to separate specular and diffuse reflection under spherical illumination. This enables independent estimation of diffuse normal maps per color channel and a specular normal map, in addition to specular intensity and diffuse albedo.
3. Rendering with independently recovered diffuse and specular normal maps, allowing detailed surface reflection and an approximation of translucent diffuse reflectance to be rendered with a local shading model.
4. A novel scanning system that uses these techniques in conjunction with structured light to recover high resolution surface geometry and aligned diffuse albedo and specular intensity maps.

2. Background and Related Work

Our techniques build on previous work in reflectance scanning, photometric stereo, and polarization-based analysis of reflected light:

Reflectance Scanning Techniques. Most reflectance acquisition techniques (e.g. [SWI, DHT*, LKG*03, GTHD, WMP*06]) use photographs of an object under different lighting conditions to acquire its reflectance properties. One of the most significant challenges is to record an object's specular reflectance behavior. Since the position of a specular highlight depends on the lighting direction and the surface normal, many images taken under different lighting directions are required to sample the appearance of the specular lobe. Even if the object is assumed to be all of the same material, many photographs are typically required [MWL*, MPBM03]. This forces most techniques to adopt one of two approaches.

The first is to take many images of the object under a dense sampling of lighting directions [DHT*, CGS, WMP*06]. This leads to data-intensive capture processes, requiring either long capture times (complicating the measurement of human subjects) or specialized high-speed cameras [WGT*05] that add expense and can compromise image quality. Furthermore, this approach remains difficult to apply to objects exhibiting mirror-like specularity; even a thousand lighting directions is insufficient to observe polished reflections without aliasing.

The second approach is to observe the specular lobe's spatial behavior across certain highlights on the object and to extrapolate this behavior across the surface of the entire object [LKG*03,ZERB]. These techniques can produce plausible reflectance properties across an object surface, but there is no guarantee of accuracy. Also, the surface orientation detail between the specular observations must be inferred using only information from the object's diffuse reflectance, and such techniques cannot capture high-frequency surface detail of scattering materials [CGS].

Photometric Stereo and Extensions. [Woo80] introduced photometric stereo, a process of determining surface orientation of a Lambertian surface from its appearance under multiple illumination directions using a simple linear system. Extensions of photometric stereo have allowed orientation to be recovered in the presence of interreflections [NIK91] as well as specularity [Sch93, MZKB]; a recent survey of photometric stereo work in this area is reviewed in [BP03].

Another line of work has endeavored to estimate surface orientation specifically from its specular reflection. Like our work, many of these use extended lighting patterns in order to reliably observe the specular reflection for a certain range of angles. [Ike81] used light from fluorescent tube lights reflecting onto diffuse surfaces to estimate surface orientation assuming mirror-like specular reflection. [HBKM] used the eye's specular reflection of a pattern of rings and spline-based surface fitting to model the shape of the cornea.

More recently, [CGS] have combined specular normals obtained using many lighting conditions with positional information to create highly detailed geometry. A key goal of our work is to obtain similarly detailed geometry for a wider angular range of specular normals using only a few lighting conditions.

Environment matting [ZWCS] measures the direction of rays from the background reflecting from or refracting through an object, making use of several extended illumination patterns created by a video monitor. [CZH*] captured similar information using a single multicolored gradient pattern to enable real-time environment matte capture. [TLGS05] went a step further, actually recovering surface normals and object geometry from these reflected directions. They recovered reflected directions using four multicolored illumination patterns that offered a good compromise of rapid capture and robust recovery. We also use four extended lighting conditions; however, we make use of full sphere illumination and of polarization to capture independent specular and diffuse normals and reflectance for an entire object.

[NIK90] demonstrated the joint recovery of surface orientation and reflectance properties for objects exhibiting both diffuse and specular reflection using several extended light sources. Our goals and approach are similar, but our technique requires fewer lighting conditions and supports the capture of distinct surface normals for diffuse and specular reflection as well as distinct surface normals for different diffuse channels. [DHT*] used many directional illumination conditions and colorspace diffuse-specular separation to recover separate surface normal maps for the diffuse and specular components of faces, and observed that the surface normals derived from diffuse reflectance lacked the sharp details of the specular normals due to translucency. Our work obtains similar results but with far fewer lighting conditions, allowing for significantly higher-speed, higher-resolution capture.

Polarization-Based Analysis of Reflected Light. Analysis of object shape and reflectance can greatly benefit from being able to determine what portion of observed light arises from diffuse reflection and what portion from specular reflection. Polarization can be a useful tool for accomplishing this. Diffuse and specular reflection components differ in that specular reflection alters polarization in a deterministic manner in accordance with the Fresnel equations, while diffuse reflection produces outgoing light which is generally close to unpolarized. [Wol, Wol89] demonstrate the use of two images taken with vertical and horizontal polarizers in front of the camera to separate diffuse and specular reflection assuming that the subject consists of a homogeneous material. [NFB97] combines this approach with colorspace analysis to separate reflectance components of heterogenous materials. Recently, [CLFS] proposed a technique for scanning translucent 3D objects using polarized illumination.

When the incident illumination can also be polarized, this separation is simplified [Mer84]. [DHT*] used a pair of images, one cross polarized and one parallel polarized, to separate diffuse and specular reflection components for characterizing the reflectance of human faces. In this work we extend this linear separation capability to full spheres of incident illumination, and also demonstrate the usefulness of circularly polarized illumination fields for diffuse-specular separation.

3. Computing Normals from Gradient Illumination

In this section we examine the relationship between surface normals and reflectance function centroids measured through gradient illumination. Due to the different behavior of diffuse and specular reflectance functions with respect to the surface normal, we discuss both cases separately.

We will show that in both cases only four spherical illu-

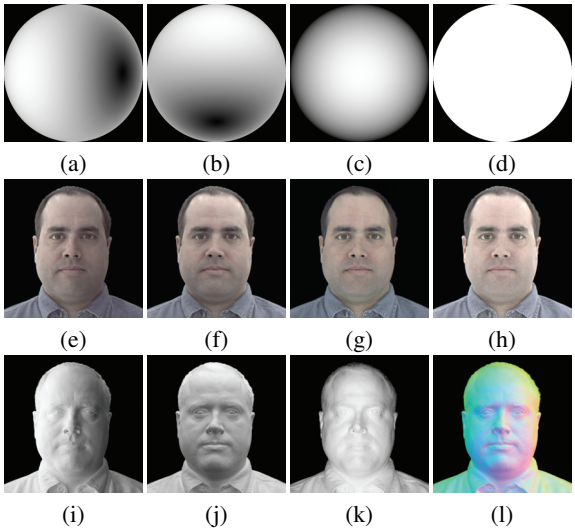

Figure 2: ***Gradient illumination conditions, ratio images, and surface normal map. (a-d)*** *The four spherical gradient lighting conditions (P'_x, P'_y, P'_z, P_c) as reflected in a mirrored sphere.* ***(e-h)*** *The gradient lighting conditions lighting a face using a spherical lighting apparatus.* ***(i-k)*** *Ratio images of (e-g) divided the uniformly lit image (h).* ***(l)*** *Surface normal map estimate derived from (i-k) using $rgb \in [0,+1]^3$ values to indicate surface normal coordinates $\vec{n} \in [-1,+1]^3$ (see online version for color).*

mination patterns $P_i(\vec{\omega})$ defined on the sphere of directions $\vec{\omega} \in \Omega$ are needed. These four patterns are: a constant pattern $P_c(\vec{\omega}) = 1$, a linear gradient along the x-coordinate $P_x(\vec{\omega}) = \omega_x$, a linear gradient along the y-coordinate $P_y(\vec{\omega}) = \omega_y$, and a linear gradient along the z-coordinate $P_z(\vec{\omega}) = \omega_z$. Note that $\vec{\omega} = [\omega_x, \omega_y, \omega_z]$ are the components of the normalized vector corresponding to the direction $\vec{\omega}$.

Lambertian Surface Reflection. Consider the Lambertian *BRDF* defined over incident illumination $\vec{\omega}$ and normal $\vec{\mathbf{n}}$ by: $R(\vec{\omega}, \vec{\mathbf{n}}) = \rho_d F(\vec{\omega}, \vec{\mathbf{n}})$, where F is the foreshortening factor $\max(\vec{\omega} \cdot \vec{\mathbf{n}}, 0)$, and ρ_d is the diffuse albedo. The observed reflectance L_i from a view direction $\vec{\mathbf{v}}$, excluding occlusions and inter-reflections, under a spherical illumination pattern P_i is:

$$L_i(\vec{\mathbf{v}}) = \int_\Omega P_i(\vec{\omega}) R(\vec{\omega}, \vec{\mathbf{n}}) d\vec{\omega}. \quad (1)$$

In case of the gradient along the x-coordinate P_x the reflectance can be derived as follows. We first express the integral in Equation (1) in a local shading frame $\mathbf{T} = [\vec{\mathbf{s}}, \vec{\mathbf{t}}, \vec{\mathbf{n}}]^T$, where $\vec{\mathbf{s}}$ and $\vec{\mathbf{t}}$ are freely chosen but orthogonal to each other and $\vec{\mathbf{n}}$, such that the normal direction $\vec{\mathbf{n}}$ is aligned to $\vec{\mathbf{z}} = [0,0,1]$, and $\vec{\omega}' = [\omega'_s, \omega'_t, \omega'_n] = \mathbf{T}\vec{\omega}$. Equation (1) becomes:

$$L_x(\vec{\mathbf{v}}) = \int_\Omega (\omega'_s s_x + \omega'_t t_x + \omega'_n n_x) R(\vec{\omega}', \vec{\mathbf{z}}) d\vec{\omega}'. \quad (2)$$

Due to the fact that $R(\vec{\omega}', \vec{\mathbf{z}})$ is not dependent on the first two components of $\vec{\omega}'$ (due to $\vec{\mathbf{z}}$), the integral over the first two terms in Equation (2) equals 0. Working out the third term gives:

$$L_x(\vec{\mathbf{v}}) = n_x \left(\frac{2\pi\rho_d}{3} \right). \quad (3)$$

A similar result can be found for the y and z gradients. In all three cases the corresponding component of the diffuse normal is scaled by the constant $\frac{2\pi\rho_d}{3}$. This constant can be divided out by normalizing the observed values.

The gradients P_x, P_y, and P_z are defined in a $[-1,+1]$ range. Because it is impossible to emit light with negative intensity, we translate and scale the intensity values to a $[0,+1]$ range: $P'_i = \frac{1}{2}(P_i + P_c)$. The reflectance values under unweighted and untranslated gradient patterns can now be easily computed by: $L_i = 2L'_i - L_c$. Although we only needed three patterns to determine the surface normal, the constant illumination pattern is needed due to this practical limitation. An example of the diffuse normal recovery process for a human subject is seen in Fig. 2.

Specular Surface Reflection. Again we start with a reflectance function $R(\vec{\omega}, \vec{\mathbf{v}}, \vec{\mathbf{n}}) = S(\vec{\mathbf{r}}, \vec{\omega}) F(\vec{\omega}, \vec{\mathbf{n}})$, where $\vec{\mathbf{r}} = 2(\vec{\mathbf{n}} \cdot \vec{\mathbf{v}})\vec{\mathbf{n}} - \vec{\mathbf{v}}$ is the specular reflected direction, and S is the specular reflectance lobe. Unlike Lambertian reflectance functions, a specular reflectance lobe S is mostly symmetrical around the perfect reflected direction. We assume that this lobe S is non-zero around a small solid angle around $\vec{\mathbf{r}}$, and that the foreshortening F is near constant in this solid angle. This assumption breaks down when (1) the extent of the lobe is large (i.e., glossy reflections), and (2) at grazing angles where the cosine changes rapidly.

We again investigate the effect of the gradient over the x-coordinate P_x. Similarly as in the Lambertian case we apply a change of variables. However, in this case we align the reflected view vector with $\vec{\mathbf{z}} = [0,0,1]$. This can be accomplished by a rotation $\mathbf{T} = [\vec{\mathbf{s}}, \vec{\mathbf{t}}, \vec{\mathbf{r}}]^T$, where $\vec{\mathbf{s}}$ and $\vec{\mathbf{t}}$ are orthogonal vectors both orthogonal to $\vec{\mathbf{r}}$.

$$L_x(\vec{\mathbf{v}}) = \int_\Omega (\omega'_s s_x + \omega'_t t_x + \omega'_r r_x) S(\vec{\mathbf{z}}, \vec{\omega}') F(\vec{\omega}', \vec{\mathbf{n}}') d\vec{\omega}', \quad (4)$$

where $\vec{\omega}' = [\omega'_s, \omega'_t, \omega'_r] = \mathbf{T}\vec{\omega}$, and $\vec{\mathbf{n}}' = \mathbf{T}\vec{\mathbf{n}}$. We made the assumption that F is approximately constant in the solid angle where S is non-zero. We denote this constant by c_F:

$$L_x(\vec{\mathbf{v}}) \approx c_F \int_\Omega (\omega'_s s_x + \omega'_t t_x + \omega'_r r_x) S(\vec{\mathbf{z}}, \vec{\omega}') d\vec{\omega}', \quad (5)$$

Again the first two terms are 0 due to the symmetry of the specular lobe S around $\vec{\mathbf{r}}$, yielding:

$$L_x(\vec{\mathbf{v}}) \approx r_x c_F \int_\Omega \omega'_z S(\vec{\mathbf{z}}, \vec{\omega}') d\vec{\omega}'. \quad (6)$$

Similar expressions can be obtained for the y and z gradients. All expressions contain the same integral, which depends on the exact definition of S, and the same constant c_F. Normalizing the vector $[L_x(\vec{\mathbf{v}}), L_y(\vec{\mathbf{v}}), L_z(\vec{\mathbf{v}})]$ yields $\vec{\mathbf{r}} = [r_x, r_y, r_z]$, which is the reflected direction of the view direction. The normalized halfway vector between $\vec{\mathbf{r}}$ and $\vec{\mathbf{v}}$ corresponds to

the specular normal direction. An example of specular normal recovery can be seen in figures 9 (c) and (d).

Discussion. The two previous derivations can be intuitively understood as follows. First observe the fact that BRDFs are largely symmetrical around the normal or reflected direction for diffuse and specular BRDFs respectively. Second, it is intuitively clear that the centroid (i.e., mean) of the BRDF lies on this symmetry axis. Thus, by measuring the centroid, the normal can be easily derived. Computing the centroid (also called the *first moment* in statistics) of 1D function can be achieved by integrating this function against a linear gradient. The presented spherical gradient functions are the extentions of a linear gradient to a spherical domain.

Real reflectance functions are rarely completely diffuse or specular, and are mostly a mix of both. Ignoring either the diffuse or specular component will result in a skewed normal. To address this, separation methods based on polarization are discussed in Sec. 4.

Until now, this discussion has ignored inter-reflections and self-shadowing. These illumination effects will change the symmetry of the reflectance functions. More specifically, this implies that the first two terms in Equations (2) and (5) are probably not zero, and thus skew the derived normal directions. Furthermore, the error on the specular normal is also dependent on how specular the BRDF is. A wide specular lobe can introduce errors due to (1) variation in F over the lobe, (2) partial occlusion of the specular lobe, and (3) off-specular reflection. As observed in [NIK90], when viewpoint is fixed and lighting direction varied, off-specular reflection is not a significant effect. This is due to the tendency of errors (1) and (3) to cancel each other.

4. Separating Reflectance Components under Spherical Illumination

We have shown that we can obtain estimates of surface normals from gradient illumination for objects whose reflectance is either diffuse or specular. Since most surfaces exhibit a combination of specular and diffuse reflectance, it is desirable to estimate normals for the object from each component independently. This section presents two polarization-based techniques for separating the diffuse and specular components.

4.1. Linear Polarization

Specular reflection is a relatively simple surface interaction arising at any interface between two media with differing indices of refraction. Denoting the index of refraction of the material by n, specular reflection from the surface is governed by the Fresnel equations:

$$r_s = \frac{\sin(\theta_t - \theta_i)}{\sin(\theta_t + \theta_i)}, \quad r_p = \frac{\tan(\theta_t - \theta_i)}{\tan(\theta_t + \theta_i)} \tag{7}$$

where r_s is the ratio of the reflected to incident electric field component perpendicular to the plane of incidence, r_p is the corresponding ratio for the parallel component, θ_i is the angle of incidence and θ_t is the refracted angle, given by $\theta_t = \arcsin(\frac{1}{n}\sin\theta_i)$.

(a)

(b)

Figure 3: ***Linear polarization reflectance separation under omnidirectional illumination. (a)*** *Diffuse component.* ***(b)*** *Specular component. Both images are taken under the top gradient P'_y. The intensity of the specular image is doubled for visualization.*

These equations imply that the polarization state of the specularly reflected light is determined by the polarization state of the incident light. In contrast, diffuse reflection arises from subsurface scattering and is almost completely unpolarized regardless of the polarization characteristics of the incident illumination. This is why the diffuse and specular components of reflected light can be effectively separated by controlling the polarization state of the incident light while also measuring the polarization state of the reflected light.

For a single camera and assuming the light source lies in a horizontal plane including the camera and subject, [DHT*] used linear polarizers placed over the light source and camera to perform this separation. This was done by placing a vertical polarizer over the light source, which given the horizontal plane of incidence assures that the outgoing specular reflection will also be vertically polarized. A horizontal polarizer in front of the camera will then block all of the specularly reflected light, but only half of the diffusely reflected light, yielding an image $I_1 = \frac{1}{2}I_D$. Similarly, a vertical polarizer in front of the camera will yield $I_2 = \frac{1}{2}I_D + I_S$. The diffuse and specular image components are then $I_D = 2I_1$ and $I_S = I_2 - I_1$. The results of diffuse-specular separation using linear polarizaiton are seen in Fig. 3 and Fig. 5.

Because we wish to illuminate our subject with many lights simultaneously and do not wish to limit the light placement to a horizontal plane, we must create a spherical direction field of linear polarization for the lights designed so that the light reflected specularly in accordance with Equation (7)

toward the camera viewpoint will be vertically polarized regardless of the angle of incidence (i.e. which light it originated from).

In practice we create this pattern by individually tuning linear polarizers placed over each light source on the sphere to minimize the observed specular reflection from a spherical test object as viewed through the camera's linear polarizer. This pattern can also be found through numerical optimization, the result of which can be seen in Fig. 4 (a).

As shown in Fig. 4 (b), the intensity of the reflected specular component from a surface element oriented along the halfway vector depends not only on the specular albedo ρ_s but also on the incident light direction itself as specified by the Fresnel equations. This effect must be divided out when interpreting the reflected specular light as ρ_s. Near the Brewster angle, the specular reflection becomes very weak, making recovery of ρ_s unreliable for these angles.

A disadvantage of separating reflectance components with linear polarization is that the required polarization direction field is highly dependent on the viewing direction. An alternate approach would be to dispense with the restriction that all specular reflection toward the camera must share the same polarization state, and instead take three images with linear polarizations spaced 60° apart. This would allow any polarization orientation to be simulated. We do not adopt this approach for two reasons. First, we found the recovery to be too ill-conditioned to permit accurate computation of the depolarized diffuse component in the presence of a strong specular component. Second, it is still necessary that the specular reflection from *any given pixel* be linearly polarized. This requires that the polarization direction field be nearly constant over the specular lobe. However, because it is impossible to specify a continuous direction field on a sphere, there must always be some illumination directions that violate this requirement. For a single viewpoint, it is possible to circumvent this problem by placing the discontinuities at the Brewster angle, as the pattern in Fig. 4 (a) does; however, this cannot be accomplished for multiple cameras at the same time. These shortcoming led us to explore circular polarization for obtaining view-independent diffuse/specular separation.

4.2. Circular Polarization

Circular polarizers have chirality but not direction. Distributing circular polarizers of the same chirality over a sphere provides a polarization field which is not only smoothly varying but also symmetric with respect to choice of camera viewpoint. This preserves the viewpoint independence established by using linear gradient patterns in the first place.

The disadvantage of this approach is that because the ratios r_s and r_p differ, circular polarization typically becomes elliptical on reflection. Thus separating diffuse and specular reflections accurately requires the capture of *four* images,

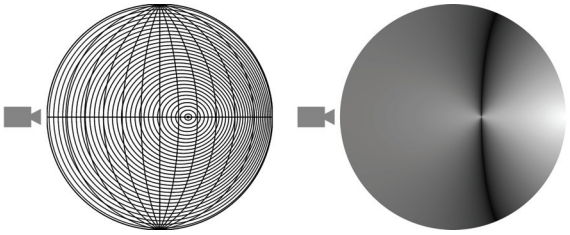

(a) (b)

Figure 4: ***(a)*** *Hemispherical angular map representation of directions of linear polarization for independently observing the specular and subsurface reflection for a full sphere of illumination. The full sphere is derived by mirroring this hemisphere. The camera is assumed to lie to the left.* ***(b)*** *Attenuation of specular reflection, represented in the same space as (a). Near the Brewster angle the reflection is weak, making recovery of true specular reflectance unstable.*

three to characterize the polarization ellipse and one to characterize the amount of unpolarized light. In addition, this recovery process is very sensitive to noise and miscalibration of the polarization filters, particularly when the specular intensity is high relative to the diffuse intensity.

Instead, we take advantage of the fact that r_s and r_p are roughly equal for moderate angles of incidence, causing the polarization of the specularly reflected light to remain nearly circular. Capturing only two images, one with a circular polarizer and another with the same filter placed backwards, then allows specular and diffuse components to be separated in the same way as for the linearly polarized case. In practice we find that the diffuse and specular are well separated for a cone of light sources surrounding each camera. Typically the half-angle of this cone is about 70°. Beyond this angle the separation degrades rapidly; however, it can be argued that surfaces more oblique than this should be captured from a different viewpoint anyway to minimize foreshortening.

To perform the separation, we placed circular polarizers over each light source on the sphere. To cancel out the specular reflections we placed a circular polarizer *of the same sense* over the camera. This is because reflection from a dielectric reverses the handedness of the circular polarization. As with separation by linear polarization, this diffuse image can be expressed as $I_1 = \frac{1}{2} I_D$. By reversing the circular polarizer over the camera, we capture another image $I_2 = \frac{1}{2}(I_D + I_S)$. (Note that this differs from the I_1 of the linearly polarized case in that only half of I_S is transmitted.) The diffuse and specular components are then $I_D = 2I_1$ and $I_S = 2(I_2 - I_1)$. The results of circular separation are seen on the right of Fig. 5. As expected, they agree well with those recovered using linear polarization for those parts of the subject not too steeply inclined away from the camera. In particular, the characteristic sharp appearance of the specular normal map and blurred appearance of the diffuse normal map are clearly evident over the central portion of the subject.

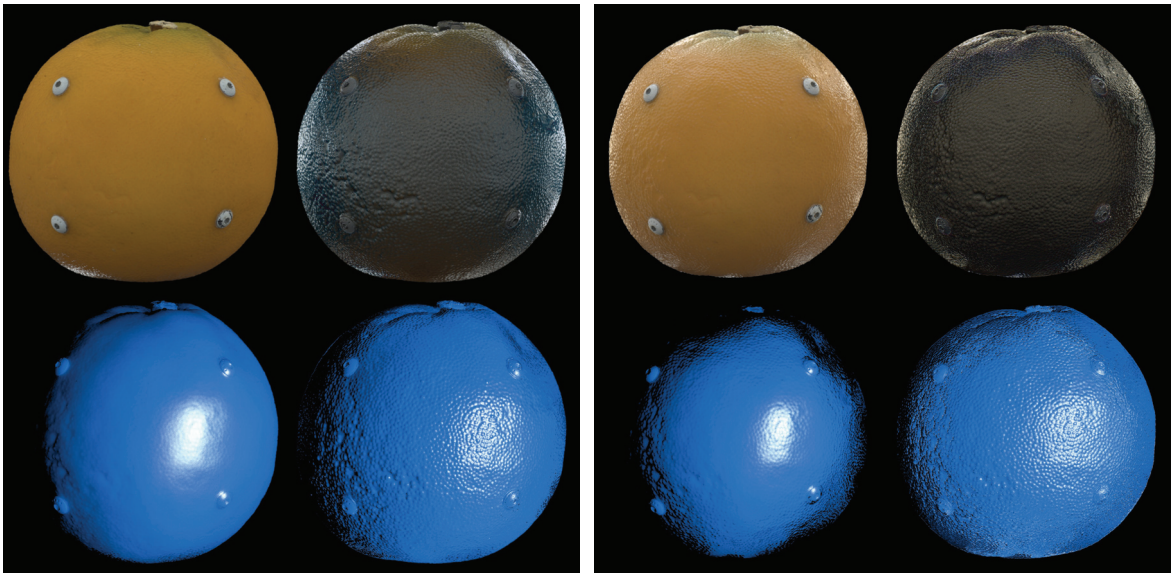

Linear Polarization | Circular Polarization

Figure 5: ***Diffuse-specular separation and normal map extraction using linear and circular polarization.*** *Shown for each of the techniques are:* ***(Upper-left)*** *Example of separated diffuse component for full sphere illumination.* ***(Upper-right)*** *Example of separated specular component for full sphere illumination.* ***(Lower-left)*** *Shaded visualization of diffuse normals.* ***(Lower-right)*** *Shaded visualization of specular normals.*

5. Results and Applications

In this section we describe four applications of the acquisition technique and show the results achieved with each. These applications include real-time capture of a performing subject, high-resolution geometry scanning, capturing and rendering with hybrid surface normals, and normal estimation of highly specular objects.

The primary lighting apparatus we use to create the gradient lighting patterns consists of 156 LED lights placed on the vertices and edges of a twice-subdivided icosahedron similar to that of [WGT*05]. For polarized patterns, individual linear or circular polarizers are placed over each light. An additional continuous-field illumination device useful for capturing highly specular objects is described in Sec. 5.4.

For all results except those described in Sec. 5.1, we use one or two digital SLR cameras. For linearly polarized illumination, a linear polarizer is mounted on a servomotor in front of the camera, allowing the polarizer to be rapidly flipped on its diagonal between horizontal and vertical orientations. For circular polarization, we manually flip a circular polarizer placed in front of the camera.

5.1. Normal Map Capture of a Live Performance

Following [WGT*05], we can acquire an animated normal map of a live performance using time-multiplexed illumination. However, since we need just four lighting conditions, we can acquire normals using a much slower frame rate than the thousands of frames per second used in [WGT*05]. In this example we use three digital video cameras running at 120fps, cycling through the four lighting conditions at 30Hz to capture the performance. This frame rate is attainable with a variety of machine vision and motion picture cameras. As in [WGT*05], we use optical flow to temporally align each set of four conditions. A principal benefit of determining surface normals using gradient illumination is that the same set of four patterns works equally well for any viewpoint, with each yielding surface normal estimates in world space. This follows from the fact that images of a subject photographed under the four gradient patterns can be transformed to any other set of gradient illumination directions using a rotation matrix [RH02]. Fig. 6(a-c) shows normal maps of a dynamic expression simultaneously captured for left, center, and right views using the same four gradient lighting conditions; the environment-mapped renderings in Fig. 6(d-f) reveal the captured surface normal detail.

5.2. High-Resolution Geometry Scanning

We can capture high-resolution surface geometry of an organic subject by combining a medium-resolution structured light scan with high-resolution surface normals based on the specular reflectance. We use a structured-light stereo technique that uses a stereo pair of SLR cameras with a video projector placed between them. The projector projects a sequence of four colored stripe patterns S_i and one uniform white pattern W. We compute correspondences between the cameras by finding best matches of the ratios $\frac{S_i}{W}$. In capturing a human subject, we first take the eight images (four gradient patterns under two linear polarization states) and then the five structured light images for a total of 13 images. Using Canon 5D cameras in "burst mode", this requires just five seconds to capture data at 12 megapixel resolution.

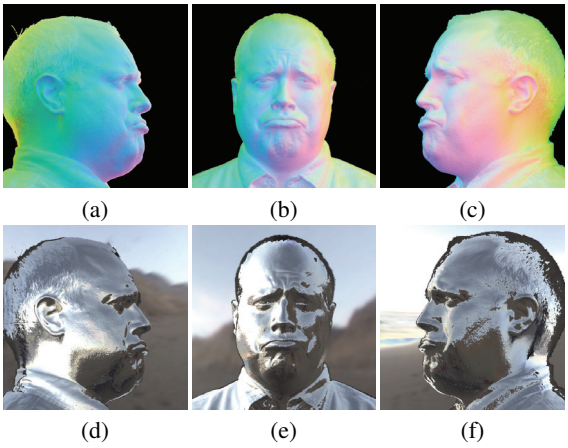

Figure 6: ***Normal map capture of live performances. (a,b,c)*** *Surface normal maps from three viewpoints derived from the same set of four gradient illumination conditions of a live performance.* ***(d,e,f)*** *Environment mapped renderings using the maps to show normal detail.*

Due to image noise and the limited resolution of the projector, the structured light scan includes some biasing and noise at high frequencies as seen in Fig. 7(a). Additional smoothing of the surface geometry results from the light scattering within the subject as noted in [CGS]. To recover the missing detail, we use the recovered diffuse or specular normals and an embossing technique. We first smooth the structured light scan using bilateral denoising [FDCO]. Next, we create a surface normal map from the smoothed mesh. We extract the high frequency details of the estimated normals using high-pass filtering and then add these details to the smoothed geometric normals. Finally, we optimize the mesh vertices to match this assembled normal map using an embossing process as in [NRDR05]. Geometry recovered in this way using diffuse or specular normals is shown in Fig. 7(b,c). As expected, the specular normals recover detailed surface shape much better than the diffuse normals.

Fig. 7 shows a comparison of the derived meshes to laser scanned geometry acquired at 0.1mm resolution from the company XYZRGB. Because the original orange object exhibits significant specularity and tranlucency, it was not suitable for laser scanning. We instead followed standard procedure and made a high-quality plaster cast for the reference laser scan. We aligned our object meshes and the reference scan using ICP, and chose a linear profile to examine. The geometric deviations in Fig. 7(e) are less than 0.1mm in depth, which is small enough so that it is difficult to determine whether they arise from our technique, from the laser scanning, or from inaccuracies in the casting. Visually, the geometry acquired using the specular normal map Fig. 7(c) appears to have the better surface detail compared to the laser scan as shown in Fig. 7(d).

Fig. 10(a) and (d) show high-resolution geometry for a face and a hand derived from structured light and specular

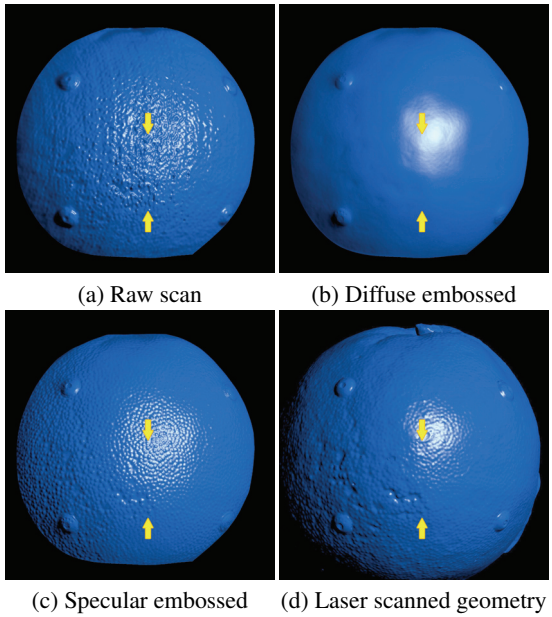

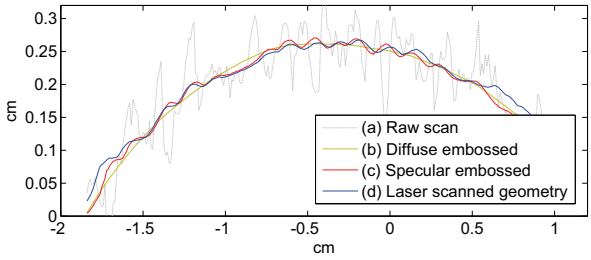

(e) Plot of depth value

Figure 7: ***High-resolution geometry capture. (a)*** *Shaded rendering of geometry recovered only using structured light.* ***(b)*** *Geometry from (a) embossed with diffuse normals derived from linear polarization in Fig. 5.* ***(c)*** *Geometry from (a) embossed with the specular normal map in Fig. 5.* ***(d)*** *Validation rendering of high-resolution (0.1mm) laser scanned geometry of a plaster cast of the orange.* ***(e)*** *Plots of a slice through the surface structure of (a-d). The yellow arrows indicate the locations of the beginning and the end of the plot samples.*

normal maps. Both of the models contain more than one million vertices. Skin details, such as wrinkles, pores, and palm prints, are well-preserved. We used these meshes, the recovered specular and diffuse albedo maps, and manually selected subsurface scattering and specular roughness parameters to generate the realistic renderings (Fig. 10(c) and (f)), using the SSS rendering technique of [JB02].

5.3. Rendering with Hybrid Normal Maps

Fig. 8 shows that for translucent surfaces the three color channels of the diffuse (subsurface) reflection and the intensity channel of the specular reflection each produce different estimates of the surface normals. Generally, the specular normals derived from the surface reflection tend to be the

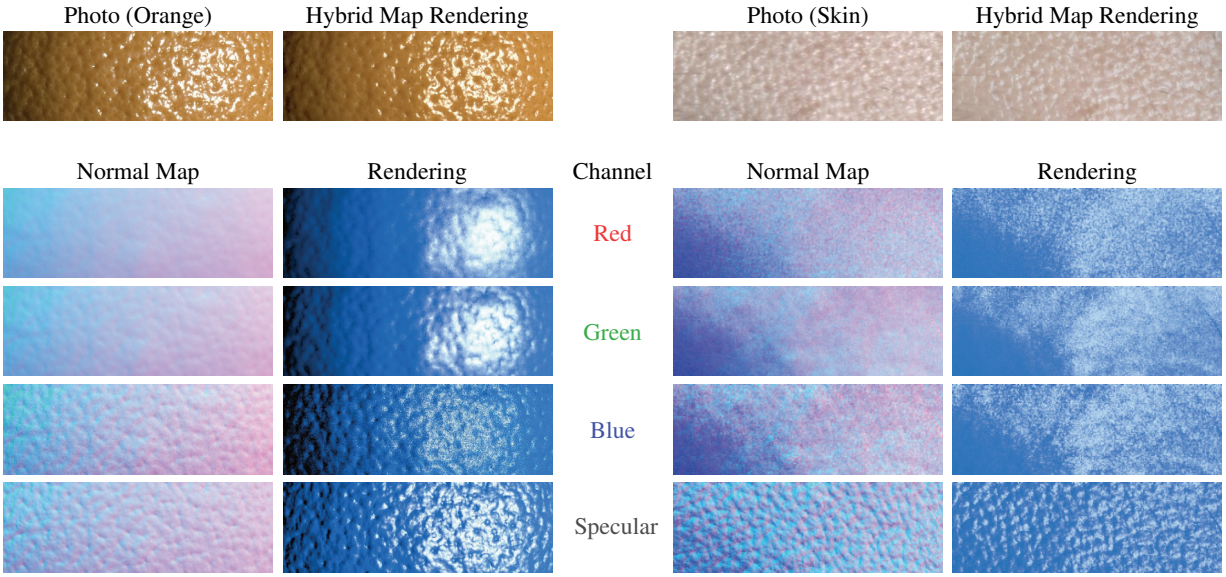

Figure 8: ***Normal map comparison for specular and diffuse color channels of an orange and human skin.*** *In both cases, shorter-wavelength illumination yields more surface normal detail. The blue channel of the orange is extremely dark leading to noisy blue normals. In both cases, the specular channel contains the most surface detail.*

best estimates of true surface orientation. The diffuse normals tend to be "bent" on a local scale due to the diffusion of light within the material, producing a normal map that appears smoother than the true surface. The three diffuse normal maps for the RGB channels show this effect to varying degrees due to the wavelength-dependence of the scattering behavior of light within the material. For skin, the red channel normals exhibit the most smoothing, and the blue channel normals exhibit the least.

We can use these different normal estimates to visually approximate some of the appearance of translucent materials in a real-time rendering framework. We do this using *hybrid normal mapping* wherein we shade the diffuse component of each color channel using the diffuse normals recovered from that channel, and similarly shade the specular component using normals recovered from the specular channel. This is easily accomplished using GPU multi-texturing. We use the images taken under the full-on pattern as the diffuse and specular albedo maps, and to create realistic renderings we manually choose a specular roughness parameter based on a reference image under point-source illumination.

Fig. 10(b) and (e) are rendered from a relatively low-resolution model using hybrid normal mapping. The results exhibit a locally translucent appearance for the face and hand while revealing surface details in the specular highlights. Fig. 10(g-i) show the difference between rendering with separately recovered surface normals for the diffuse red, green, and blue components and simply using the diffuse normal recovered from the green channel for all three color channels. Fig. 10(h) demonstrates the improved reproduction of wavelength dependent translucency by using different normals maps for each color channel. Additional hybrid map renderings shown at the top of Fig. 8 and in Fig. 1(c) can be seen to closely match actual images of the subject under point-source illumination.

5.4. Normal Map Acquisition for Highly Specular Surfaces

The discrete lighting sphere used so far cannot estimate surface normals for very shiny surfaces since most surface orientations will reflect little or no light from the sources. To apply our technique to a highly specular surface, we use a continuous-field hemispherical lighting device similar to that in [PHD06]. We use a $1m$ diameter hemisphere painted with rough silver paint to reflect light from a source at one focus onto a subject at the other focus as shown in Fig. 9(a). To generate the gradient illumination patterns, we place a 5cm frosted spherical light bulb at the source focus, 10cm to the right of center. We place a round occluder between the light and the subject to prevent light from reaching the subject directly. We view the subject through a hole at the apex of the hemisphere seen in the middle of Fig. 9(a).

With the bulb on, the dome reflects isotropic light onto the subject from nearly the complete hemisphere, allowing us to photograph the subject under uniform hemispherical illumination corresponding to pattern P_c. To generate the gradients P'_x, P'_y, P'_z, we use a cap to cover the half of the light bulb facing the $-x$, $-y$, and $-z$ axes, respectively. Each position generates a linear gradient of irradiance on the surface of the dome since the visible proportion of the light bulb varies with the cosine of the angle from which it is viewed, and this cosine is equal to the x, y, or z coordinate of the corre-

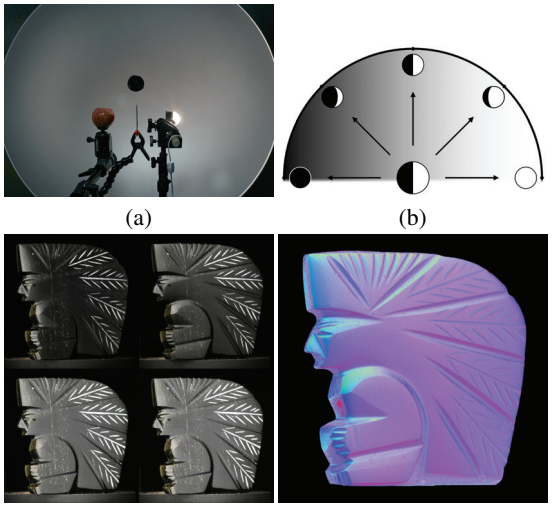

Figure 9: ***Acquiring normals for a highly specular surface. (a)*** *Reflective hemisphere to create continuous gradient illumination, showing the viewing aperture, subject, and half-covered light bulb.* ***(b)*** *Diagram of a half-covered light bulb producing gradient illumination.* ***(c)*** *Left, top, front, and full gradient illumination conditions on a polished obsidian sculpture.* ***(d)*** *Computed normal map from the specular reflection.*

sponding gradient axis (see Fig. 9(b)). The dome reflects this gradient pattern back onto the subject.

Fig. 9(c) shows a polished obsidian sculpture photographed in this manner under the four spherical lighting conditions. Since the object's diffuse component is black, the preponderance of the reflected light is from specular reflection. By processing these images as described in the specular discussion of Sec. 3, we obtain the normal map shown in Fig. 9(d). Since the specular reflections are sharp, the lack of illumination from the opposite hemisphere affects only the normals greater than 45° away from the camera. The recovered normal map and its associated specular intensity map is used to relight the sculpture under animated point-source lighting in the accompanying video.

6. Future Work

Our results suggest a number of avenues for future work. Currently, our technique does not allow for the estimation of specular roughness; we instead choose the specular lobe shape manually for our renderings. It is possible that using higher-order gradients of illumination would allow the estimation of the specular lobe width at each surface point. Alternatively, it could be fruitful to combine our approach with that of [BMA] and reflectance sharing [ZERB], using a few point-light directions (perhaps lit simultaneously) to produce observable specular highlights across the surface. These observations could be extrapolated across the entire face, benefitting from the detailed specular surface normal information that our approach records from the gradients.

We have so far only used the technique to record diffuse normal maps for a live-action performance, since our current setup does not support capturing both polarization states at video frame rates. To acquire live specular normal maps from linear or circular polarization, we could either use ferroelectric shutters to rapidly alternate between the two states, or we could film the states simultaneously with two cameras through a beam splitter.

7. Conclusion

We have presented a novel method for acquiring specular and diffuse normal maps using a small number of gradient illumination patterns. Due to the low number of patterns, the method is suitable for real-time normal capture using time-multiplexed illumination. We presented two polarization techniques for diffuse-specular separation under spherical illumination, allowing high-resolution surface detail to be recorded from the specular component of hybrid surfaces. Using this we demonstrated a high-resolution scanning system that uses these techniques in conjunction with structured light to produce fast high-resolution scans of organic materials. Finally, we demonstrated the use of different recovered normal maps for shading the diffuse and specular components to efficiently approximate some of the effects of surface translucency.

Acknowledgements

We would like to thank Mark Bolas from the USC School of Cinematic Arts for building the motor-driven polarizer flippers, Helmut Kungl of XYZRGB, Inc. for scanning the cast of the orange, and Per Einarsson, Andrew Gardner, Akihiro Fujita, TOPPAN Printing Co. Ltd., Bill Swartout, Scott Fisher, Randy Hill, Randolph Hall, and Max Nikias for their support and assistance with this work. This work was sponsored by the University of Southern California Office of the Provost and the U.S. Army Research, Development, and Engineering Command (RDECOM). The content of the information does not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.

References

[BMA] BABA M., MUKUNOKI M., ASADA N.: Estimating roughness parameters of an object s surface from real images. In *SIGGRAPH 2004 Posters Proceedings*.

[BP03] BARSKY S., PETROU M.: The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE Trans. Pattern Anal. Machine Intell. 25*, 10 (2003), 1239–1252.

[CGS] CHEN T., GOESELE M., SEIDEL H.-P.: Mesostructure from specularity. In *CVPR '06: Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2, pp. 1825–1832.

[CLFS] CHEN T., LENSCH H. P. A., FUCHS C., SEIDEL H.-P.: Polarization and phase-shifting for 3D scanning of translucent

objects. In *CVPR '07: Proceedings of the 2007 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.

[CZH*] CHUANG Y.-Y., ZONGKER D. E., HINDORFF J., CURLESS B., SALESIN D. H., SZELISKI R.: Environment matting extensions: Towards higher accuracy and real-time capture. In *Proceedings of SIGGRAPH 2000*, pp. 121–130.

[DHT*] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. In *Proceedings of SIGGRAPH 2000*, pp. 145–156.

[FDCO] FLEISHMAN S., DRORI I., COHEN-OR D.: Bilateral mesh denoising. In *Proceedings of SIGGRAPH 2003*, pp. 950–953.

[GTHD] GARDNER A., TCHOU C., HAWKINS T., DEBEVEC P.: Linear light source reflectometry. In *Proceedings of SIGGRAPH 2003*, pp. 335–342.

[HBKM] HALSTEAD M., BARSKY B. A., KLEIN S., MANDELL R.: Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals. In *Proceedings of SIGGRAPH 96*, pp. 335–342.

[Ike81] IKEUCHI K.: Determining surface orientation of specular surfaces by using the photometric stereo method. *IEEE Trans. Pattern Anal. Machine Intell. 3*, 6 (1981), 661–669.

[JB02] JENSEN H. W., BUHLER J.: A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics 21*, 3 (2002), 576–581.

[Lan02] LANDIS H.: Production-ready global illumination. Course Notes for SIGGRAPH 2002 Course 16, RenderMan in Production, 2002.

[LKG*03] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics 22*, 2 (2003), 234–257.

[Mer84] MERSCH S.: Polarized lighting for machine vision applications. In *Proceedings of RI/SME Third Annual Applied Machine Vision Conference* (1984), pp. 40–54.

[MGW] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *Proceedings of SIGGRAPH 2001*, pp. 519–528.

[MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Transactions on Graphics 22*, 3 (2003), 759–769.

[MWL*] MARSCHNER S. R., WESTIN S. H., LAFORTUNE E. P. F., TORRANCE K. E., GREENBERG D. P.: Image-based BRDF measurement including human skin. *Eurographics Rendering Workshop 1999*.

[MZKB] MALLICK S. P., ZICKLER T. E., KRIEGMAN D. J., BELHUMEUR P. N.: Beyond Lambert: Reconstructing specular surfaces using color. In *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2.

[NFB97] NAYAR S., FANG X., BOULT T.: Separation of reflection components using color and polarization. *International Journal of Computer Vision 21*, 3 (1997), 163–186.

[NIK90] NAYAR S. K., IKEUCHI K., KANADE T.: Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Journal of Robotics and Automation 6*, 4 (1990), 418–431.

[NIK91] NAYAR S. K., IKEUCHI K., KANADE T.: Shape from interreflections. *International Journal of Computer Vision 6*, 3 (1991), 173–195.

[NRDR05] NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHI R.: Efficiently combining positions and normals for precise 3d geometry. *ACM Transactions on Graphics 24*, 3 (2005), 536–543.

[PHD06] PEERS P., HAWKINS T., DEBEVEC P.: A reflective light stage. *ICT Technical Report ICT-TR-04.2006* (December 2006).

[RH02] RAMAMOORTHI R., HANRAHAN P.: Frequency space environment map rendering. *ACM Transactions on Graphics 21*, 3 (July 2002), 517–526.

[Sch93] SCHLÜNS K.: Photometric stereo for non-lambertian surfaces using color information. In *CAIP '93: Proceedings of the 5th International Conference on Computer Analysis of Images and Patterns* (1993), pp. 444–451.

[SWI] SATO Y., WHEELER M. D., IKEUCHI K.: Object shape and reflectance modeling from observation. In *Proceedings of SIGGRAPH 1997*, pp. 379–388.

[TLGS05] TARINI M., LENSCH H. P. A., GOESELE M., SEIDEL H.-P.: 3D acquisition of mirroring objects using striped patterns. *Graphical Models 67*, 4 (2005), 233–259.

[WGT*05] WENGER A., GARDNER A., TCHOU C., UNGER J., HAWKINS T., DEBEVEC P.: Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics 24*, 3 (2005), 756–764.

[WMP*06] WEYRICH T., MATUSIK W., PFISTER H., BICKEL B., DONNER C., TU C., MCANDLESS J., LEE J., NGAN A., JENSEN H. W., GROSS M.: Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics 25*, 3 (2006), 1013–1024.

[Wol] WOLFF L.: Using polarization to separate reflection components. In *In CVPR '89: Proceedings of the 1989 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 363–369.

[Wol89] WOLFF L.: Material classification and separation of reflection components using polarization radiometric information. In *DARPA89* (1989), pp. 232–244.

[Woo80] WOODHAM R. J.: Photometric method for determining surface orientation from multiple images. *Optical Engineering 19*, 1 (1980), 139–144.

[ZERB] ZICKLER T., ENRIQUE S., RAMAMOORTHI R., BELHUMEUR P.: Reflectance sharing: Image-based rendering from a sparse set of images. In *Proceedings of Eurographics Symposium on Rendering 2005*, pp. 253–264.

[ZWCS] ZONGKER D. E., WERNER D. M., CURLESS B., SALESIN D. H.: Environment matting and compositing. In *Proceedings of SIGGRAPH 1999*, pp. 205–214.

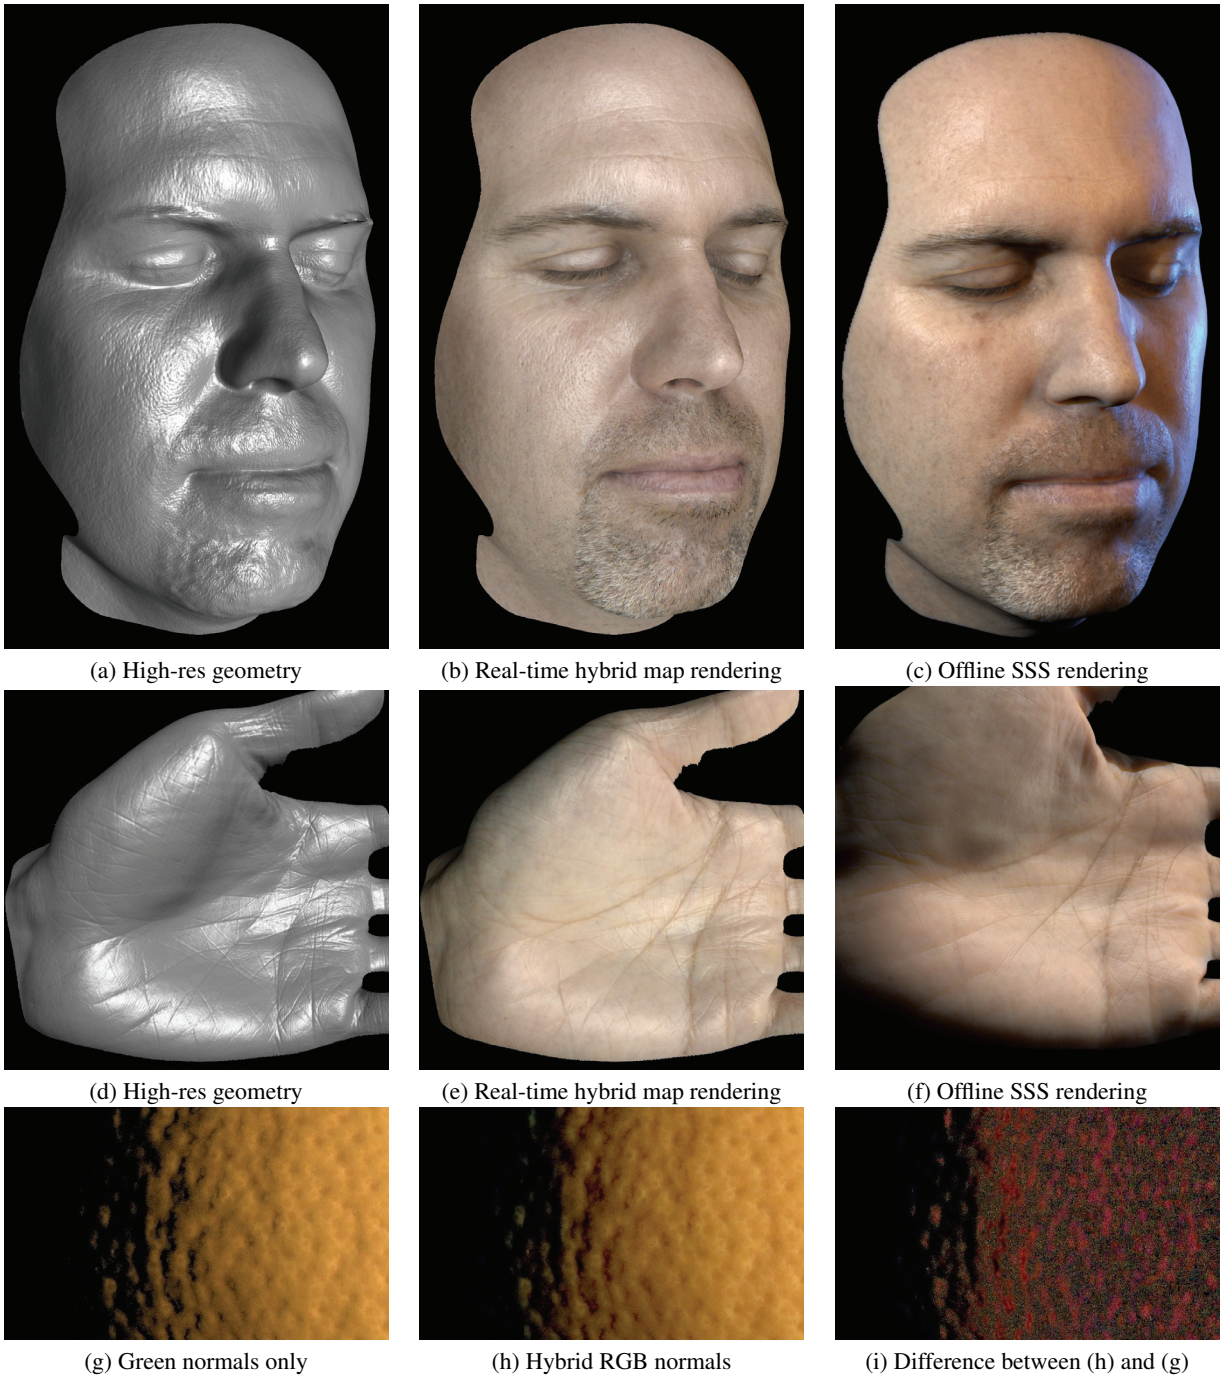

(a) High-res geometry (b) Real-time hybrid map rendering (c) Offline SSS rendering

(d) High-res geometry (e) Real-time hybrid map rendering (f) Offline SSS rendering

(g) Green normals only (h) Hybrid RGB normals (i) Difference between (h) and (g)

Figure 10: ***Additional Results.*** ***(a)*** *High-resolution face geometry obtained by embossing a specular normal map onto a structured-light 3D scan.* ***(b)*** *Real-time rendering of low-resolution face geometry with hybrid normal mapping for diffuse and specular reflectance, with a manually chosen specular roughness.* ***(c)*** *Offline rendering using high-resolution geometry, estimated diffuse albedo and specular intensity, and global illumination with subsurface scattering.* ***(d)*** *High-resolution hand model obtained by the same method of (a).* ***(e)*** *Real-time hybrid normal rendering of the hand.* ***(f)*** *Offline rendering of the hand using high-resolution geometry and SSS scattering.* ***(g)*** *A detail from the orange in Fig. 1, rendered using just the green channel diffuse normal map to shade all three diffuse color channels.* ***(h)*** *Detail from Fig. 1(c), rendered using independently recovered RGB normal maps to shade the RGB channels.* ***(i)*** *Difference image (at +2 stops) between (h) and (g), showing the appearance of additional reddish light within the concavities when hybrid normal rendering is used.*